

METHODS OF CONTROLLING TRANSMISSION POWER LEVELS IN AIR INTERFACE CHANNELS

Field of the Invention

The present invention relates to wireless communications, and in particular although not exclusively, to networks of indoor collocated or overlapping wireless coverage areas especially those intended for operation in unlicensed spectrum such as the Industrial Scientific and Medical (ISM) bands.

Background of the Invention

Collocated wireless coverage areas using different channels can still interfere with each other, adjacent channel suppression depends largely on the receiver implementation and transmitted spectral mask. Similarly, overlapping or adjacent wireless LAN coverage areas using uncoordinated spread spectrum mechanisms such as direct sequence or frequency hop spread spectrum also “intentionally” interfere with each other to a certain extent when the same or adjacent channels are used at the same time. Typically these channels are defined by frequency, (as used in frequency division multiple access) time slot (time division multiple access), code (code division multiple access) or a combination. The level of interference on a given channel will depend on many factors including the signal strength of the wanted and interfering signals which in turn will depend on their transmission power level and the signal propagation environment, for example whether there is likely to be a strong multipath fading of the wanted or interfering signal as is typically difficult to predict in indoor environments.

The level of interference which is acceptable can also depend on what type of data is being transmitted. For example non-real time data such as email can be re-transmitted if

not received correctly. However other types of real time data such as video base layer frames are more critical and sensitive to delays such as re-transmissions and therefore require a higher quality channel to deliver them in good time.

“Dynamic Reconfiguration and Efficient Resource Allocation for Indoor Broadband Wireless Networks” by Tim Farnham and Brian Foxon, Universal Personal Communications, 1998, ICUPC '98, IEEE 1998 International Conference, Vol 1, 1998, pages 53-57; describes an indoor radio access network having clusters of access points or base stations that use a combination of time division and frequency division multiple access mechanisms. The network sets different interference thresholds for type 1 (video) and type 2 (data) traffic, the type 1 threshold being lower as video is delay sensitive which means it needs to be sent first and accurately to avoid the need for re-transmission. The network switches between different channels when too many type 1 packets are waiting for transmission through an existing channel because its interference level is too high. The trigger for a channel change is based on the average buffer occupancy and error rate over time.

Summary of the Invention

In general terms, the present invention provides a method of controlling a wireless network having at least one base station and one or more mobile stations each containing one or more radio transceivers, in which the power levels of channels allocated between a base station and a mobile station are dynamically varied in order to reduce the interference level on channels which are carrying traffic with a high quality of service level (QoS) requirement.

By controlling channel power levels, the network can optimise conditions for channels transmitting traffic (such as packets) requiring a high QoS. This may involve reducing the power levels on channels carrying traffic with a lower QoS requirement, which may, but need not, result in higher interference levels or larger delays for that lower QoS traffic. This enables more critical or sensitive data to be transmitted with greater

accuracy, by providing a low interference channel when required, but without building into the network the need to maintain these low interference channels the whole time which would reduce capacity. Less critical or sensitive data can be transmitted using higher interference or lower QoS channels, thereby increasing the capacity of the network. If need be this data can be retransmitted as it is less sensitive to delays.

In particular, in one aspect the present invention provides a method of controlling the transmission power associated with an air interface channel, the method comprising determining a quality of service measure for a signal to be transmitted on a second air interface channel; determining an interference measure corresponding to the interference on said second air interface channel which will be caused by transmissions on said first air interface channel; requesting the changing of the transmission power associated with the first air interface channel when the signal is to be transmitted on the second air interface channel, when the quality service measure is above a predetermined value and when the interference measure is above a predetermined value for said quality of service value.

Preferably the transmission power associated with the first air interface channel is changed when the signal is to be transmitted on the second air interface channel.

Alternatively, if this is not possible, the transmission associated with the first air interface channel may be postponed or suppressed when the signal is to be transmitted on the second air interface channel.

The quality of service measure and the interference measure determinations may be based on the prediction of likely quality of service requirements and interference levels.

Determining the interference measure may determine whether the signal to be transmitted or received over the first air interface channel will occur within the same time frame as transmissions on the second air interface channel.

The air interface channels are associated with different air interface technologies, such as WLAN IEEE802.11 standards, or Bluetooth™ for PANs.

In another aspect the present invention provides a method of allocating power levels to transceivers in a wireless network having a base station and two or more mobile stations using said transceivers; the method comprising: determining whether a signal to be transmitted over a first said transceiver has a high QoS indicator; determining whether any of said other transceivers will cause a predetermined level of interference to occur on said first transceiver when said signal is to be transmitted; instructing the reduction of the transmission power levels of one or more said other transceivers when said traffic signal is transmitted over said first transceiver.

Preferably the transceivers are collocated.

In yet another aspect the present invention provides a method of allocating power levels to transceivers in a wireless network having a base station and two or more mobile stations using said transceivers; the method comprising: allocating one said transceiver for high quality of service signals; instructing another said transceiver to reduce its transmission power level if it would otherwise interfere with a high quality of service signal by above a predetermined amount.

In other aspects the invention provides corresponding apparatus for implementing these methods, as well as software or computer or processor code which is capable of controlling a computer or processor to carry out these methods.

Brief Description of the Drawings

In order to better understand the invention, embodiments will now be described with reference to the following drawings, by way of example only and without intending to be limiting, in which:

Figure 1 is a schematic of a wireless network;

Figure 2a is a schematic of a base station for the network of figure 1;

Figure 2b is a schematic of a mobile station for the network of figure 1;

Figure 3 is a flow chart showing a method of operating the base or mobile stations of figure 2;

Figure 4a is a schematic of the transmission management arrangement for the base station of figure 2a;

Figure 4b is a schematic of the transmission management arrangement for the base station of figure 2b;

Figure 5 is a graph of the performance of a network of collocated Bluetooth™ piconets;

Figure 6 is a schematic of a Bluetooth™ access point; and

Figure 7 is a schematic of a device supporting both Bluetooth™ and IEEE802.11b protocols.

Detailed Description

Figure 1 shows an embodiment comprising a wireless network having one or more base stations or hubs 2 wirelessly coupled to a number of mobile stations 1. The network is arranged into one or more coverage areas 5. These coverage areas 5 might be PAN's (Personal Area Network) each comprising a number of collocated Bluetooth™ piconets each servicing a different mobile station 1 within the PAN. Alternatively the network might be a WLAN (Wireless Local Area Network) comprising adjacent or overlapping IEEE802.11 coverage areas 5 or a combination of PAN and WLAN. Wireless signals 3

couple each mobile station 1 to its corresponding base station 2. These signals 3 may also travel into adjacent areas 5 causing (especially) co-channel and adjacent channel interference (shown generally as 4).

Figure 2a is a schematic of a base station or access point 2a providing four air interface channel processing chains (shown generally as 16a, 16b, 16c, 16d) for communicating with mobile stations 1 within its coverage area 5. Each chain 16a – 16d has associated transmission and/or reception functions for a radio access technology (such as radio signal modulation and demodulation, synchronisation, start of frame detection, frequency hopping, medium access control, link management, packet reception and transmission) or transceivers 14a – 14d, each having a transmission power control function 15a – 15d which controls the level of transmission power for that transceiver's signal 3 and synthesisers 17a, 17b, 17c, 17d for setting the transmit and receive frequency (or channel) of the transceiver 14.

The base station 2 may also comprise a multiplexer/demultiplexer 11 for interfacing between the separate air interface chains 16 and a combined data link 10 to the outside world; for example phone line, Ethernet LAN or other base stations in the case of a WLAN.

Each air interface processing chain 16 is also associated with a transmission management function (TM) 12a – 12d, which performs intelligent cooperative power management operations for the corresponding transceivers 14a, 14b, 14c, 14d and controls the transmission power levels of each transceiver 14 in the coverage area (eg BluetoothTM piconet 5) through the corresponding power control functions 15, which can override air interface specific power control algorithms implemented within the transceivers 14. The detailed functionality required within the transmission management function 12 to support this power management can be seen with references to figures 3 and 4.

The transmission management (TM) function 12a for a corresponding transceiver 14a first determines whether the next (or subsequent) packet(s) to be transmitted has a high

Quality of Service requirement (QoS). This is indicated by a predetermined field within the packet or by analysing the packet type and size or by fixed priorities assigned to the different transceivers 14 and the method of packet classification implemented within the demultiplexer. The indication may be binary, that is either high (critical) or low (best effort); or there may be various levels of QoS requiring different handling.

If the next or subsequent packet(s) requires a high QoS (i.e. higher than can be achieved without intervention), then functionality within the TM function 12a determines estimates of the level of interference on the channel 3 that the transceiver 14a will be transmitting or receiving on for its next (or a predetermined) transmission slot (or other transmission period over which the packet will be transmitted) by interrogating the other TM functions 12b, 12c, 12d about their next transmission slots (in particular the frequency or channel and transmit / receive timing). By determining whether the transceivers 14b, 14c or 14d will be transmitting in the next slot, the TM function 12a can estimate the interference this is likely to cause in its receiver (14a). This is also the level of interference that is likely to be observed at the receiving mobiles 1 of the transmission from the transmitter 14a. This is described in more detail later.

The present TM function 12a can also be configured to interrogate the mobile stations 1 of its own coverage area 5 (eg piconet) for their intended transmissions over the next transmission slot. It is also possible that the AIP function can estimate which mobiles are (or are likely) to be transmitting in the next time period. For example by analysing the transmission from the base station to the mobile or previous transmissions from the mobile to the base station.

In the case of a WLAN, in addition to determining whether the other 14b, 14c, 14d transceivers of the present base station 2a will be transmitting when present transceiver 14a intends to transmit or receive the high QoS packet(s), the TM function 12a can also determine whether radio transceivers in adjacent areas 5 will be transmitting during the next slot. It does this by interrogating the other base stations 2b, etc associated with those other areas 5 over the data link 10, which also provides the traffic and other control signals. The present TM function 12a may just interrogate the base stations 2a,

etc of adjacent areas 5, or it may interrogate base stations 2 which are perhaps two or more areas 5 distant from the present base station 2a, depending on the signal propagation environment. Similarly, the mobile stations 1 communicating with these other base stations 2 may also be interrogated, or their likely transmissions estimated.

From all this information, or a sub-set of it, the TM function 12a is able to estimate the interference which will or is likely to occur in its associated channel of the transceiver 14a and the mobiles connected to this transceiver over the next transmission / reception slot. If this is above a predetermined threshold for a high QoS packet, then the TM function 12a selects an interfering transceiver 14b (or transceivers) and requests that the power on this be reduced by a predetermined amount for the next transmission slot in order to reduce the interference on the air interface processing chains 16a associated with the present transceiver 14a (and the mobiles connected via 14a) which is to transmit (or receive) the high QoS packet(s). Where there are a series of incoming or outgoing high QoS packets for transceiver 14a, the other transceivers 14b, 14c, 14d may simply be requested to reduce transmission power for the whole period of this series.

The transceiver 14b that is requested to reduce power receives the “reduce Tx power” request and determines whether the packet can still be transmitted at reduced power (with increased probability of packet loss). If not, then the packet is suppressed until the next transmission opportunity. If there is conflict due to too many high priority packets that cannot be suppressed that are being scheduled simultaneously, the TM can make a request to change channels, which may or may not be possible depending on the radio access technology. Otherwise the associated power control function 15b is instructed to reduce power by the requested amount for the next transmission slot (or period).

Whilst this method has been described in relation to the next transmission slot for ease of explanation, other arrangements could be used for example using transmission (and reception) sequences (as described in “An asynchronous time division multiplexing-multiple access protocol for indoor wireless multi-service networks”, by T Farnham, Universal Personal Communications 1997, ICUPC '97 IEEE 6th International Conference, 12-16 Oct 1997).

Figure 2b is a schematic of a mobile station 1 having one air interface processing chain 16ma for communicating with the base station 2a of figure 2a. The mobile station 1 has an associated transmission and/or reception function or transceiver 14ma which has a transmission power control function 15ma which controls the level of transmission power for that transceiver 14ma and a synthesizer function 17ma that sets the transmit and receive frequency (or channel) of the transceiver. A transmission from the mobile to the base station can be received by any one of the transceivers 14a, 14b, 14c and 14d in the base station. Likewise the transmission from the base station to the mobile can be from any transceiver 14a, 14b, 14c and 14d. However, the channels (frequency) that transceivers 14a and 14ma are tuned to (via the synthesisers 17a and 17ma) and air interface technology (modulation and coding schemes etc.) used by the transceivers 14a and 14ma must be the same for the transmission to be successfully received.

The mobile station 1 has an AIP function 12ma which can receive requests to raise or lower transmitter power by the base station transceiver (14a). Normally this would be performed with air interface specific power control commands (i.e. directly between 14a and 14ma). Therefore, existing methods can be utilised for communicating power control commands to mobile stations. It could also utilise the same commands in the reverse direction (for requests from the mobile station to the base station transceiver 14). However, to convey request covering longer time periods with more advanced power management commands (such as with priority and QoS information) would require separate signalling methods (for example transmission sequences can be multicast or broadcast using user data or control channels).

The mobile may also contain more than one transceiver and in this case it could be implemented in an identical manner to the base station 2.

Figures 4a and 4b are schematics showing in more detail the TM functions 12 for the base stations 2 and mobile stations 1 respectively. For ease of explanation, reference is only made to other TM's and transceivers' existing within a single coverage area 5, that is by a single base station 2 and mobile stations 1 within the coverage area 5, as is

typically the case in a PAN. It being understood that this can be expanded to incorporate TM's and transceivers associated with base stations and mobile stations in other coverage areas 5 such as in a WLAN for example, with the use of appropriate control signals over the network.

Referring to figure 4a and 2a, the TM function 12a of base station 2a is associated with a transceiver 14a used for both reception and transmission, and comprises a first buffer 22 coupled between the demultiplexer 11 and the transmission/reception circuitry (shown as separate transmitter 14at and receiver 14ar functions for ease of explanation). The packet classification performed within the demultiplexer can utilise combinations of source / destination addressing, packet type, flow labels or packet QoS indicators to determine which transceivers to route the packets destined for the mobiles.

The power management functionality can be distributed between transceivers (14a, 14b, 14c and 14d) with one TM 12 per transceiver 14, as shown in figure 2a, alternatively a single TM can be implemented to control all transceivers. In either case the detailed functionality of the TM is the same as shown in figure 4a.

A scheduling function 24a, controls the transmission packet queue to determine when a packet requiring high QoS is present, and the transmission time this is scheduled to be transmitted. The required QoS is typically given in a predetermined field in each packet, or can be based on packet type, packet size, or destination/source address or packet flow label.

A suppression function 23a either instructs the power control function (15at) within the transceiver to reduce power level of a transmission or suspends the intended transmission of a packet until the next transmission opportunity (following the suppression period). The suppression can be achieved in a number of ways. For instance by reducing power level to in effect zero, which is likely to result in lost packets or higher layer retransmissions. Alternatively, the transceiver could be instructed to enter a sleep mode for a predetermined length of time.

A performance and traffic monitoring function 27a determines performance parameters associated with the receiver 14ar when receiving signals, for example measurements can be made of signal strength, bit error rate, packet latency, packet failure or retransmission rate or any combination of these. The monitoring function (27a) also observes traffic patterns or levels of activity, which are not necessary for normal operation, but can be used to make predictions about likely levels of interference by analysing packet arrivals and packet sizes. These various measurements are assessed by a performance assessment function 25a which monitors the affect of interference from other transceivers over time. This information can be used by the scheduler 24a to determine whether a packet with high QoS can be sent given transceiver 14a interference levels.

A performance prediction function 26a estimates the transceivers (14a) likely interference levels for upcoming transmission slots based on output from the performance assessment function 25a as well as other data from other TM functions. This other data might include future transmission options (taking into account frequency hop and channel offset patterns transmission deadlines and permissible power levels). It may even include information on what QoS level (or priority) the packets they will be transmitting in particular slots are intended to have.

From these predictions, the scheduler 24a determines whether a high QoS packet can be transmitted during a predetermined transmission slot based on the predicted interference level for the channel at that time. If the interference level is too high, the scheduler 24a selects a transceiver 14b (or transceivers) to reduce power and issues a corresponding request via the suppression function 23b to for that transceiver (14b). The request may include the priority level of the packet, the source(s) that needs to be suppressed, the length of time that suppression is required, and the level of suppression (for example in dB). If the scheduler's request is successful then it transmits the high QoS packet at the allotted transmission time slot. If the scheduler's 24a request is not successful, it either selects another transceiver 14c to reduce power (which it has determined will cause some interference and if this would satisfy the requirements) and performs another request, or delays the high QoS packet to another time slot. Alternatively the TM can be configured to instruct a change of channel (i.e. carrier frequency) which may or may

not be possible depending on the air interface technology. A prediction is then made about the likely interference for that time slot and if necessary a reduce power message(s) is sent to a selected transceiver(s).

The scheduler 24a also receives requests to suppress a particular transmission slots from the associated suppression function 23 when a request to suppress (or reduce transmit power) from another transceiver (16b) is successful (i.e. the suppression function has accepted to reduce power or suppress the transmission).

The receiving part 14ar of the transceiver 14a is coupled to the multiplexer 11 of the base station 2a, which forwards incoming packets to the appropriate destination.

Figure 4b shows a schematic of the TM function 12am for a mobile station 1 which is analogous to that of the base station, except that control signals such as reduce power requests are all sent over a wireless link.

A given PAN or WLAN (or combination) may incorporate both base station(s) and mobile stations implementing these functions, alternatively only one of these may use this arrangement.

The above described functions may be implemented by software and/or hardware accelerators such as ASIC's or FPGA for example.

The above described functions may be implemented by software and/or hardware accelerators such as ASIC's or FPGA for example.

An example implementation is described in more detail below. The base station 2 has four collocated transceivers 14a, 14b, 14c and 14d. The first two transceivers supporting the (cck modulated) 802.11b (14a and 14b) air interface and the second two the (frequency hopping fsk modulated) BluetoothTM air interface (14c and 14d). In this example the QoS requirements can be met with fixed priority assignment in the order of 16a (being highest priority) to 16d (being lowest priority). Each transceiver has its own

power control functionality 15a, 15b, 15c, 15d respectively. For simplicity only one mobile is associated with each transceiver 14a, 14b, 14c and 14d (denoted 1a, 1b, 1c and 1d respectively). The mobile stations can receive TM specific power control commands but cannot perform other more complex power management functions.

The packet classification within the demultiplexer 11 ensures that the packets are routed to the different transceivers (and on to the mobiles 1a, 1b, 1c and 1d) based on destination address and packet type. The specification of the timing of suppression requires a shared clock of sufficient resolution to provide accurate estimation of when interference will be caused. In this case a clock with resolution of one BluetoothTM time slot (625microseconds) is used. Also, the power control method must be precise enough to allow for the different levels of performance required. The transceivers 14 each support 4 power levels (corresponding to approximate transmit power levels of 20dBm, 10dB, 4dBm, 0dBm). The performance assessment and prediction method utilise packet success rate (using packet CRC check) and latency.

When a packet is waiting transmission at the transceiver 14a, the corresponding scheduler function (24a) uses the performance predictor (26a) to discover whether there will be acceptable interference or not, using timing information obtained from the other transceivers scheduling functions (24b, 24c and 24d), frequency channel offset pattern information and the current transceivers performance assessment function (25a), which provides an indication of past performance. If the interference is deemed to be unacceptable (for example if the centre frequency between channels is less than a predetermined amount Δf^1_{ir}) a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 1 level (i.e. 1 level below that of the 14a transceiver) for the period in which the performance predictor (26a) indicates that interference will be caused. If the separation between channel centre frequencies is less than Δf^2_{ir} then a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 2 levels (i.e. 2 levels below that of the 14a transceiver). If the separation between channel centre frequencies is less than Δf^3_{ir} then a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 3 levels (i.e. 3

levels below that of the 14a transceiver). If the separation between channel centre frequencies is less than Δf_{tr}^4 then a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 4 levels (i.e. 4 levels below that of the 14a transceiver).

If the transmit power level of the current transceiver 14a is less than the highest level, the scheduler (24a) can simultaneously request to increase the power level of transceiver 14a (to the power controlling function 15a) while also sending the suppression requests to the other transceivers suppression functions, (23b, 23c and 23d).

Likewise, the reception of a packet at the transceiver 14a is predicted by the monitoring function (27a) observing the traffic pattern. The pattern in this case is a regular group of packets at a fixed interval and so the arrival of the next group of packets is easily predicted. When the next group of packets is expected, the performance monitor function (27a) uses the performance predictor (26a) to discover whether there will be acceptable interference or not, using timing information obtained from the other transceivers scheduling functions (24b, 24c and 24d), frequency channel offset pattern information and the current transceivers performance assessment function (25a), which provides an indication of past performance. In the same way as for packet transmission, if the interference is deemed to be unacceptable (in this case if the centre frequency between channels is less than Δf_{rt}^1) a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 1 level (i.e. 1 level below that of the 16a transceiver) for the period in which the performance predictor (26a) indicates that interference will be caused. If the separation between channel centre frequencies is less than Δf_{rt}^2 then a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 2 levels (i.e. 2 levels below that of the 14a transceiver). If the separation between channel centre frequencies is less than Δf_{rt}^3 then a request is made to the appropriate transceivers suppression functions (23b, 23c and 23d) to suppress transmission by 3 levels (i.e. 3 levels below that of the 16a transceiver). If the separation between channel centre frequencies is less than Δf_{rt}^4 then a request is made to the appropriate transceivers

suppression functions (23b, 23c and 23d) to suppress transmission by 4 levels (i.e. 4 levels below that of the 16a transceiver).

The same functionality is implemented in the other TM's 12b, 12c and 12d.

In this case frequency offset has been used as the main decision criteria. Alternatively, in fixed frequency systems for example, the modulation scheme and/or antenna direction (in the case of directional antennas) could be used.

The embodiment is applicable to any radio networking protocol, for example BluetoothTM for a PAN arrangement, and HiperLan or IEEE802.11 for a WLAN arrangement. The channels may be defined by any known means, for example by frequency, code, time slot or spatial multiplexing. It could also be applied in device-to-device communications where two or more channels are used simultaneously.

The embodiment efficiently supports different levels of quality of service with bursty traffic characteristics. This enables more efficient utilisation of radio resources and improves support for different services and applications. Figure 5 shows the performance of BluetoothTM with collocated piconets. The packet failure rate could easily be unacceptable for many applications. Therefore suppression of interfering piconets when interference is unacceptable makes this traffic supportable (in this example the packet failure rate would fall from over 10% with 3 collocated piconets to around 3%). Other services can easily sustain 10% packet failure rate with standard retransmission mechanisms.

Figure 6 shows a BluetoothTM access point implementation. In this arrangement the piconets are overlapping in physical space, each providing a single channel to wirelessly couple a single mobile station to the access point.

One or more of these piconets is designated to support higher QoS than the others, for example to carry MPEG4 base layer frames. Expected interference is thus measured or

predicted with respect to the dedicated high QoS piconets, the other piconets having their power levels suppressed as required.

Figure 7 shows a device which supports both BluetoothTM and IEEE802.11 protocols. Here BluetoothTM might support a voice call whereas 802.11 supports data transfer. The QoS of the BluetoothTM connection is more important than the 802.11 connection, as packet retransmission is not a good option for a voice call. Therefore the performance assessor suppresses the 802.11 transmissions while the BluetoothTM system is receiving or transmitting. In this case the performance prediction is simply based on frequency separation between the BluetoothTM and 802.11 systems.

The invention has been described with respect to embodiments thereof. Alterations and modifications as would be obvious to those skilled in the art are intended to be incorporated within the scope thereof.